STRUCTURAL ECONOMETRIC ANALYSIS OF RANDOMIZED CLINICAL TRIALS: THE CASE OF ACTG175

Tat Y. Chan and Barton H. Hamilton
Olin School of Business
Washington University in St. Louis
Campus Box 1133
One Brookings Drive
St. Louis, MO 63130
USA

chan@olin.wustl.edu; hamiltonb@olin.wustl.edu

April 2002

VERY PRELIMINARY AND INCOMPLETE

PLEASE DO NOT CITE OR QUOTE

1. INTRODUCTION

The "gold standard" approach in the literature for the evaluation of a treatment on an outcome has been to conduct a randomized clinical trial (RCT) in which patients are randomly assigned to treatment and control groups. The basic framework for assessing the causal impact of a treatment involves the construction of the *potential outcome* (Rubin (1974, 1978), Heckman and Robb (1985)) associated with each level of the treatment. Treatment effects are then given by differences in potential outcomes for various treatment levels. Because outcomes are only observed for one level of the treatment at a particular time, the investigator is faced with the task of constructing what the outcome might have been had the subject received another level of the treatment. In a clinical trial, treatment effects are calculated by comparing the outcomes of the treated with those of the controls. Randomization addresses the possibility that bias may be introduced into the estimation if individuals were able to choose treatment status on their own. Typically, the treatment effect is then given by the regression coefficient on the treatment variable in an Ordinary Least Squares (OLS) regression of the outcome of interest on the binary treatment indicator as well as other covariates.

The standard approach for evaluating treatment effectiveness in a clinical trial may lead to biased results in longitudinal trial settings due to attrition, which, in many cases, changes a randomized clinical trial into an observational study if such dropouts are non-random across the treatment arms of the trial. While the biased induced by dropouts has been increasingly recognized as problem in evaluating clinical trials, it has generally been viewed as a statistical issue that should be controlled for in one fashion or another (Scharfstein et al. (1999)). However, recent work by Philipson and Desimone (1997) and Philipson and Hedges (1998) argues that dropout behavior provides information that is useful in the evaluation of the treatment. In particular, the decision to

leave the RCT depends not only on the impact of the trial drugs on the "publicly observed" measures of health status, such as the CD4 count, that is the typical focus of treatment efficacy evaluation, but also upon "private" information observed by the patient, such as the pain or discomfort the he may feel when taking the treatment medication. We label these latter conditions as *side effects* since they typically reflect short-term discomfort associated with consumption of the drug. Philipson and Desimone (1997) argue that participants engage in "subject sampling" in which they attempt to learn about the direct and side effects of the treatment, since they have a strong incentive to do so. Consequently, dropout behavior provides insight not only into the direct effectiveness of the treatment, but also potential side effects that may not be easily measured or are privately observed by the subject, as well as treatment options that lay outside the clinical trial.

This paper constructs structural models of RCT subject behavior in which trial participants decide whether to drop out by comparing the utility generated by remaining in the trial, which is a function of both direct and side effects of the treatment received, with the returns obtained by seeking care outside the trial. We estimate two alternative models, related to specifications found in the literature on demand for pharmaceuticals, that impose different assumptions on the behavior of subjects and on how information is revealed in the RCT: (1) the *learning model* takes the approach of Ching (2000) and Currie and Park (2000) and assumes that subjects maximize current period utility, but are initially uncertain of the effectiveness of the treatment over time and infer this over time via Bayesian updating (although the side-effects of the treatments are observed immediately by the subject); (2) in the *forward-looking learning model*, subjects learn about treatment effectiveness over time, but decide whether to remain in the trial for one more period by comparing the expected present value of associated with remaining in the trial with that obtained by leaving and receiving the outside option. A similar model is used by Crawford and Shum (2000) to analyze the demand for anti-ulcer medications in a non-RCT setting.

We use our framework to analyze data from the AIDS randomized clinical trial ACTG 175 (see Hammer et. al. (1996)). ACTG 175 was a randomized double blind clinical trial to evaluate the effectiveness of combination therapy versus monotherapy for HIV infected individuals with CD4 cell counts of between 200 and 500/mm^3. CD4 counts are a widely used marker for the status of an individual's immune system, and hence the progression of AIDS in the patient. CD4 counts were measured at trial baseline and then at weeks 8, 20, 32…104 of the trial. The potential outcomes for this paper are the subjects' CD4 count at each trial week under AZT+ddC and each week's CD4 count under AZT. A notable feature of ACTG 175 was that roughly half of the subjects dropped out by the second year of the trial. While the substantial attrition was noted in the initial evaluation of the trial (Hammer et. al. (1996)), only Scharfstein et al (1999)) has attempted to account for dropout in the estimation of the treatment effect in the trial. However, Scharfstein et al analyze a cross-section of data from the trial and do not consider the evolution of attrition over the course of the trial.

The first objective of our paper is to assess the impact of the combination therapy zidovudine plus zalcitabine (AZT+ddC) relative to zidovudine (AZT) on CD4 counts, accounting for attrition. Recent studies by Chickering and Pearl (1997) and Heckerman and Shachter (1995) emphasize the importance of accounting for heterogeneity when constructing treatment effects in health care studies. We therefore allow the impact of trial treatments on CD4 counts to vary across trial subjects. The second objective of the paper is to examine the extent to which treatments differ in their side effects. Randomization at baseline provides information through which this latter set of treatment effects is identified. If attrition is greater in one trial arm than another, conditional on effectiveness, then the side effects of the particular treatment are likely to be greater, since unobserved patient characteristics are balanced at baseline.

While we have not completed the estimation of the forward-looking learning model, our

initial findings show that combination therapy increases CD4 counts by about 8.7 percentage points relative to monotherapy after 8 weeks in the trial. This effect increases to 31.9 percentage points by week 104. However, the positive treatment effect of AZT+ddC is offset by the fact that it appears to have significantly greater side effects than AZT alone. The combination of these two opposing effects helps to explain why dropout is initially higher for combination therapy patients, although by the end of the trial attrition is greater for subjects receiving monotherapy. As one might expect, the learning model estimates suggest that trial participants have optimistic prior beliefs concerning trial treatment effects. This somewhat offsets the impact of the lower observed CD4 counts of AZT patients on attrition early on in ACTG 175.

The remainder of the paper first presents initial summary evidence on the importance of accounting for attrition when estimating the impact of treatments on the evolution of CD4 counts. Section 3 then describes the structural models of dropout behaviour, while Section 4 describes the econometric implementation of the economic models. The parameter estimates are presented and treatment effects constructed in Section 5. A short conclusion summarizes the findings and suggests avenues for future research.

2. THE ACTG 175 DATA

The data for this paper comes from Aids Clinical Trial Group Study 175 (ACTG 175) that compares the impact of monotherapy to combination therapy for 2467 HIV-infected adults with screening CD4 cell counts from 200 to 500 per cubic millimeter. Subjects were recruited from 43 AIDS Clinical Trials Units and 9 National Hemophilia Foundation sites in the United States and Puerto Rico. A total of 98 trial sites were used in the trial. Individuals were randomly assigned to one of four daily treatment regimens: 400 mg of didanosine (ddI); 600 mg of zidovudine (AZT); 600 mg of zidovudine plus 400 mg of didanosine (AZT+ddI); or 600 mg of zidovudine plus 2.25 mg of zalcitabine (AZT+ddC). During 1991 and 1992, when the trial

was initiated, AZT was the standard treatment for HIV. For each subject, treatment response was measured by the longitudinal progression of CD4 counts recorded at different intervals. Subject CD4 counts were recorded baseline (week 0), week 8, and then every twelve weeks thereafter for a period of 104 weeks or more. Full details of the trial may be found in Hammer et al. (1996).

For simplicity, this analysis focuses on the treatment effect for two arms of the trial, AZT and AZT+ddC. The analysis sample consists of 1072 subjects treated at 89 different trial sites, 536 receiving AZT, 536 receiving AZT+ddC. We examine the effect of these treatments on CD4 counts for the first two years after baseline, so that a subject potentially has 10 measures of CD4 counts, recorded at weeks $w = 0,8,20...104$. The two-year window was chosen because the clinical endpoint of the trial for all subjects occurred at least two years after baseline.

Following Boscardin et al. (1998), the outcome measure used in this study is the change in the ln(CD4) count between week w and baseline week 0. This specification allows us to determine whether the rise in CD4 counts often observed 8 weeks after baseline continues to persist into subsequent weeks. Figure 1 plots the relationship between the average change in log CD4 counts and weeks in the trial for subjects in each of the two treatment arms. The CD4 profiles shown in the figure indicate that the average CD4 count of individuals receiving AZT+ddC actually increases in week 8 relative to baseline. While mean CD4 counts decline over the remainder of the two-year period, the average remains near the baseline for this group. In contrast, subjects receiving AZT alone experience a 5% fall in CD4 counts by week 8. By week 104, CD4 counts are roughly 15% below baseline.

Table 1 presents estimates from a regression of the change in ln(CD4) counts relative to baseline on indicator variables for the week of the trial and interactions between trial week and receipt of AZT+ddC. Specification (1) shows that the differences in CD4 profiles across

treatment arms shown in Figure 1 are statistically significant. For example, the change in CD4 count is 10% higher for those receiving AZT+ddC in week 8, relative to those taking AZT. By week 104 of the trial, this difference increases to approximately 19%. Specification (2) presents estimates for a specification that also includes controls for age, gender, whether the subject is white, the subject's screening CD4 count, whether the subject received ZDV prior to the start of the clinical trial, and whether the subject is HIV symptomatic. As might be expected from the random assignment of trial participants, the estimated treatment effect does not change when these controls are included in the model.

2.1 INITIAL EVIDENCE CONCERNING THE IMPACT OF ATTRITION

For the purpose of our analysis, attrition occurs when a subject ends the treatment assigned at baseline prematurely (i.e., before reaching week 104). Figure 2 plots the survivor function over the two-year period for the AZT+ddC and AZT sub-samples. Attrition appears to be a potentially important confounder in assessing the progression of CD4 counts in ACTG 175, since only about 45%-50% of subjects continue in the trial through week 104. Moreover, the figure shows "cross-over" behavior in attrition: although subjects were randomly assigned to treatment, AZT+ddC subjects are initially more likely to leave ACTG 175, but over time the attrition rate for AZT patients overtakes that for those receiving combination therapy.

The reason a subject drops out of the trial is reported in the ACTG 175 data. While the reported reasons are hardly definitive, and may mask multiple causes, some insight on the potential role of side effects in inducing attrition may be gained by examining the reason for dropout data in Table 2. The first row of the table shows that, conditional upon leaving the study for any reason, the death of a subject while participating in the trial is relatively uncommon. Moreover, few patients are explicitly removed from the trial by the ACTG 175 investigators. Consequently, dropout appears to be in large part a subject decision. Comparison of columns (1)

and (4) show that patients receiving combination therapy are more likely to request to leave the trial at some point due to toxic reactions to the treatment than are those receiving AZT alone, perhaps suggesting greater side effects of AZT+ddC. The fraction of combination therapy patients who report dropping out due to toxic reactions declines substantially for those leaving the study in weeks 8 – 44 (column (5)) is much greater than that for dropouts in weeks 56 plus (column (6)). Thus, it appears that patients learn very quickly about the side-effects of AZT+ddC. In contrast, the table shows that subjects receiving AZT are more likely to request to leave the trial without a particular reason being reported, particularly in the earlier weeks of the trial. This may reflect the weaker treatment effect of this drug. Of course, it may be the case that a subject with a strong reaction to the treatment may simply not return to the trial physician and be classified as lost to follow-up rather than leaving due to toxicity.

To provide an initial assessment of the potential impact of attrition on treatment response, profiles of the change in log CD4 counts of non-dropouts were compared with those of dropouts. Small cell sizes prevent us from constructing separate profiles for individuals dropping out in week 20, 32,...,104. Instead, subjects were classified into 3 groups: (a) those who did not drop out between week 0 and 104 (non-dropouts, 47% of subjects); (b) subjects dropping out between weeks 68 and 104 (Year 2 dropouts, 18% of subjects); and (c) subjects that dropped out between weeks 8 and 56 (Year 1 dropouts, 35% of subjects).

The impact of attrition in assessing AIDS progression in the trial is highlighted by Figures 3 and 4, which plot the mean change in ln(CD4) counts for the three groups by treatment status. Note that for Year 1 dropouts, CD4 counts are available for all subjects in week 8, but the mean value shown in the figure for weeks 20, 32, and 44 only reflect the CD4 counts of subjects surviving that long (all Year 1 dropouts have left by week 56). Similarly, for the Year 2 dropouts, CD4 counts are available for all subjects prior to week 68. Comparison of the plots in

Figure 3 for the AZT+ddC sub-sample indicates substantial differences in CD4 profiles by attrition group. The change in CD4 counts between week 8 and baseline is higher for non-dropouts than for those dropping out in Year 1. By week 44, Year 1 dropouts experience a 20% decline in CD4 counts, compared to a 5% decline for Year 2 dropouts and a 3% increase for non-dropouts. By week 92, the mean CD4 count for Year 2 dropouts is about 13% below baseline, compared with a 3% increase for non-dropouts.

While all groups in the AZT sub-sample shown in Figure 4 experience greater declines than those reported for AZT+ddC patients, the pattern across attrition groups is similar. For example, the decline in CD4 counts for Year 1 dropouts is so rapid that the change in CD4 counts in week 44 for this group is similar to the change in week 92 for Year 2 dropouts. These plots suggest that attrition is a potentially serious problem for understanding the impact of therapies on disease progression. This is confirmed by the regression estimates presented in Column (3) in Table 1. In this regression, indicator variables for the week of dropout are added to specification (2). Because the week of dropout variables refer to future events, they capture the impact of heterogeneity associated with dropout. Comparison of the estimated coefficients on the week dummies in columns (2) and (3) show that the decline in CD4 counts indicated by specification (2) is contaminated by attrition bias. As is clear from Figures 3 and 4, those remaining in the sample through week 104 have starkly different CD4 profiles compared to individuals dropping out earlier.

3. MODELLING SUBJECT BEHAVIOR IN ACTG 175

Trial subjects decide each period whether to remain in the trial or drop out and seek alternative treatment. The decision to remain in the trial depends in part on the subject's evaluation of the direct impact on health status of the treatment received in ACTG 175, denoted by H_{it}, as well as the side effects experienced by the patient when taking the trial medication, S_{it}. While H_{it} is a

measured outcome in the trial (e.g., the CD4 count) and is typically the focus of the evaluation of the efficacy of alternative treatments by trial investigators, side effects are assumed to be private information to the subject. Variation in the side effects associated with particular treatments potentially leads to situations where a treatment may have a strong positive impact on health status, but subjects receiving the treatment are much more likely to drop out of the RCT in spite of the observed impact on H_{it}.

As noted in the introduction, we consider two alternative models of attrition that differ in the restrictions imposed on subject behaviour as well as the complexity of econometric implementation. Each of the models is structural in the sense that we specify subject utility functions and separately identify direct effectiveness and side effect distributions. In the learning model, we assume that patients maximize current period utility but are not forward-looking. Subjects also immediately observe side effects S_{it}, although they do not know which arm of the trial they are in. This specification is consistent with the results shown in Table 2, which indicate that combination therapy patients quickly discover whether they have a toxic reaction to their treatment. However, the direct effectiveness of the treatment on H_{it} is not immediately known. The subject is uncertain about treatment effectiveness both because he is blinded to his assigned treatment arm, and because there is subject-level variation in the impact of the treatment on health status. Subject i learns about the effectiveness of treatment by observing the sequence $\{H_{it}\}$ over the course of the trial. This specification of utility maximization and learning is similar to the pharmaceutical demand model of Currie and Park (2000) in which beliefs concerning drug effectiveness are updated via a Bayesian learning process. Dynamic behaviour reflects the learning process as well as unexpected period-specific shocks and changes in the outside option over the course of the RCT.

We also consider a forward-looking learning model in which subjects choose to remain in the trial for one more period if the discounted stream of expected current and future utilities is greater

than the value of the outside option. The model is similar to that of Crawford and Shum (2000), although those authors did not have an observable measure of H_{it}. In this case, dynamic behaviour reflects expectations regarding future utility as well as learning.

3.1 THE SUBJECT'S DECISION PROBLEM

We now describe the forward-looking learning model and note the restrictions that are imposed by the simple learning models. To fix ideas, let $t = 0,1,\ldots,\tau$ index the time periods in the trial (corresponding to weeks 0, 8, 20,…) and $i = 1,\ldots,N$ denote subjects. The indicator variables d_{it}^{k} indicate whether the subject remains in the trial ($k = T$) or chooses the outside option and drops out ($k = o$). A subject may be in only one state in each period, so that $d_{it}^{T} + d_{it}^{o} = 1$. The objective of the subject is to choose a sequence of actions that maximizes the present value of lifetime utility:

$$(1) \quad \max_{\{d_{it}^{k}\}} E[\sum_{s=t}^{\tau} \beta^{s-t} \sum_{k} E[U_{is}^{k} \mid I_{is}] d_{it}^{k} \mid I_{it}] + \beta^{s-\tau} V_{i}^{o}(I_{it}, \tau+1),$$

where U_{it}^{k} is the period-specific flow of to the subject from alternative k, I_{it} is the subject's information set at time t, and β is the discount rate. We assume that at the end of the trial period, all subjects leave the trial and receive the discounted lifetime flow of expected utility from the outside option, $V_{i}^{o}(I_{it}, \tau+1)$.

In ACTG 175 subjects are not allowed to re-enter the trial once they drop out, so we impose the constraint $d_{it+1}^{o} = 1$ if $d_{it}^{o} = 1$. Therefore, at any time t during the trial, the lifetime flow of expected utility available to the subject if he leaves the trial is given by

$$(2) \quad V_{i}^{o}(I_{it}, t) = E[\sum_{s=t}^{\tau} \beta^{s-t} E[U_{is}^{o} \mid I_{it}] + \beta^{s-\tau} V_{i}^{o}(I_{it}, \tau+1)].$$

Given the subject's information set at time t, I_{it}, the expected return of remaining in the trial in period $t < \tau$ may be written as

$$(3)\quad V_i^T(I_{it},t)=E[U_{it}^T \mid I_{it}]+\beta E[\max\{V_i^T(I_{it+1},t+1),V_i^o(I_{it+1},t+1)\} \mid I_{it}],\quad 0\le t<\tau,$$

In the terminal period of the trial, $t = \tau$, the subject receives

$$(4)\quad V_i^T(I_{i\tau},\tau)=E[U_{i\tau}^T \mid I_{i\tau}]+\beta E[V_i^o(I_{i\tau+1},\tau+1) \mid I_{i\tau}]$$

if he remains in the trial, since once the trial is completed we assume that the patient receives the outside treatment option. The maximal lifetime utility of the subject at time t is then given by the maximum of the value of remaining in the trial for at least one more period and the value of leaving the trial and receiving the outside treatment forever, since subjects cannot re-enter the trial once they have dropped out:

$$(5)\quad V_i(I_{it},t)=\max\{V_i^T(I_{it},t),V_i^o(I_{it},t)\}.$$

For a positive discount rate, the decision-making framework outlined in equations (1) – (5) implies that subjects may remain in the trial despite low current period utility, perhaps resulting from painful side effects, if they expect the future benefits of trial participation to be particularly high. In addition, remaining in the trial for an additional period augments the subject's information set, which also may increase future utility as the patient learns more about the effectiveness of the treatment. Note that both the learning model specification assumes that $\beta = 0$ so that subjects maximize current period utility only.

3.2 SPECIFICATION OF PREFERENCES

Subject i's per-period utility obtained by participating in ACTG175 is a function of his longer term health status in period t as measured by CD4 count, H_{it}, as well as the shorter term side effects (e.g., pain or nausea associated with consumption of treatment drugs) experienced by the individual when taking the trial drugs, S_{it}. To simplify the analysis, we assume that utility is additively separable in H_{it} and S_{it}. In addition, the impact of long term health status on utility is allowed to be non-linear. Currently, we assume that utility is a linear function of side effects,

implying that

$$(6) \quad U_{it}^{T}(H_{it}, S_{it}) = H_{it}^{\gamma} - S_{it}.$$

No information is available from ACTG175 concerning the subject after he drops out of the clinical trial. Consequently, we specify the stream of lifetime utility of subjects dropping out of the trial at time t defined in equation (2) to be a linear function of observed characteristics x_{oit} as well as a stochastic component:

$$(7) \quad V_{i}^{o}(I_{it}, t) = x_{oit}\delta + \varepsilon_{iot}.$$

3.3 SPECIFICATION OF HEALTH STATUS AND SIDE EFFECTS

Let h_{it} denote the natural logarithm of health status of subject i in period t, $\ln(H_{it})$. Studies in the literature evaluating the impact of alternative treatments on CD4 counts typically use a log specification to capture the percentage impact of the treatment (Boscardin et. al. (1998)). Therefore, h_{it} is assumed to be a function of the subject's (log of) health status at the beginning of the trial, h_{i0}; the subject level response to the treatment received in the RCT; a time trend; and a period specific error term:

$$(8) \quad h_{irt} - h_{i0} = \theta_{ri} + \lambda_{r}t + \upsilon_{it},$$

where the subscript $r = 0, 1$ indicates the treatment arm the subject is randomised into. In ACTG175, $r = 1$ implies the subject receives AZT+ddC, while $r = 0$ implies receipt of AZT alone. The subject-level random effect in equation (8) is normally distributed and allowed to depend on time-invariant subject characteristics

$$\theta_{ri} \sim N(z_{i}\beta_{r}, \sigma_{\theta r}^{2}).$$

While the ACTG 175 data contain an observable measure of H_{it} (the subject's CD4 count), no direct measures of S_{it} are available. Consequently, side effects are specified to be a linear function of observed characteristics, including treatment arm, and a period-specific

stochastic component:

$$(9) \quad S_{it} = x_{Si}\alpha_i + \varepsilon_{iSt}, \qquad \alpha_i \sim N(\alpha, \sigma_\alpha^2).$$

Equation (9) incorporates a random coefficient specification to allow the impact of side effects to vary across subjects.

3.4 INCORPORATING SUBJECT LEARNING

In the models that incorporate learning, subjects decide whether to remain in the trial in period *t* before health status h_{it} is observed, and must form expectations concerning its value. Subjects do not observe the idiosyncratic component υ_{it}, but they do know its distribution. In addition, subjects are blinded as to which treatment arm they have been assigned. Consequently, from the subject's point of view at time *t*

$$(10) \quad h_{it} \sim N(h_{i0} + \mu_i, \sigma_\upsilon^2),$$

where μ_i is the unknown individual specific impact on health status of remaining in the trial. We assume that at trial baseline, subjects have common prior beliefs concerning μ_i:

$$(11) \quad \mu_i \sim N(\mu_0, \sigma_{\mu 0}^2).$$

The prior variance σ_μ^2 reflects the precision of the prior beliefs of the subjects in the trial. The prior parameters μ_0 and $\sigma_{\mu 0}^2$ are estimated as part of the econometric model.

While all trial participants have the same prior mean and variance at baseline, upon the commencement of the trial the subject observes a sequence of health status measures h_{it} , which is used to update the subject's prior beliefs according to the Bayesian rule (DeGroot (1970)). The posterior mean and variance of μ_i in period *t* is given by

$$(12)\quad E[\mu_i \mid I_{it}] = \mu_i^t = \sigma^2_{\mu,t}\left[\frac{\mu_0}{\sigma^2_{\mu 0}} + \frac{1}{\sigma^2_{\upsilon}}\sum_{s=1}^{t-1}(h_{is} - h_{i0})\right],$$

$$Var(\mu_i \mid I_{it}) = \sigma^2_{\mu,t} = \frac{1}{\frac{1}{\sigma^2_{\mu 0}} + \frac{t-1}{\sigma^2_{\upsilon}}}.$$

Similar expressions may be constructed to describe the learning process for side effects. However, in the current version of the paper we assume that side effects are learned immediately upon initiation of the trial by the patient. As a result, the idiosyncratic factors influencing side effects at time t, ε_{iSt} are observed by the subject (but not by the econometrician) when the dropout decision is made. This may be a reasonable assumption, given that many side effects reflect short-term discomfort associated with consumption of the treatment. Such features of the trial treatment seem likely to be immediately observed.

Given the utility specification in equation (7) and the assumptions regarding subject learning described above, and noting that $H_{it}^{\gamma} = \exp(\gamma * h_{it})$, the expected current period t utility of remaining in the trial given the patient's information set I_{it} may be written as

$$(13)\quad E[U_{it}^T(H_{it}, S_{it}) \mid I_{it}] = \exp\left(\gamma * (h_{i0} + \mu_i^t) + \frac{\gamma^2}{2}(\sigma^2_{\upsilon} + \sigma^2_{\mu,t})\right) - x_{Si}\alpha_i - \varepsilon_{iSt},$$

since h_{it} is a normal random variable. Equation (13) may then be substituted into equations (3) and (4) to obtain the expected return to remaining in the RCT for at least one more period.

4. ECONOMETRIC IMPLEMENTATION

This section presents the econometric method used to estimate the economic model described in Section 3. The estimation procedure is complicated by the fact that dropout decision is correlated over time, as are the observed health status measures $\{H_{it}\}$. Therefore, our econometric approach relies on simulation-based methods and proceeds in two steps. To begin, consider a subject who drops out in period t. The likelihood contribution of this subject is then

the joint probability of observing the sequence of (log) health outcomes h_{i1}, h_{i2},…,h_{it-1} and dropout in period t and is given by

$$(14)\quad L_i = \Pr(h_{i1},...,h_{it-1}, d_{i1}^T = 1,...,d_{it-1}^T = 1, d_{it}^T = 0 \mid \Theta_1, \Theta_2, x_i, z_i),$$

where the parameters from the dropout decision and health status equations are denoted by $\Theta_1 = \{\beta, \gamma, \delta, \alpha, \sigma_\alpha, \mu_0, \sigma_\mu\}$ and $\Theta_2 = \{\beta_r, \sigma_{\beta r}, \lambda_r, \sigma_\upsilon\}$, respectively, and $x_i = \{x_{Si}, x_{oit}\}$. Equation (14) can be rewritten as

$$(15)\quad L_i = f(h_{i1},...,h_{it-1} \mid d_{i1}^T = 1,...,d_{it-1}^T = 1, d_{it}^T = 0, \Theta_1, \Theta_2, x_i, z_i) * \Pr(d_{i1}^T = 1,...,d_{it-1}^T = 1, d_{it}^T = 0 \mid \Theta_1, \Theta_2, x_i, z_i).$$

Using equations (3), (4), and (7), this is equivalent to

$$(16)\quad L_i = f(h_{i1},...,h_{it-1} \mid V_{i1}^T > V_{i1}^o,...,V_{it-1}^T > V_{it-1}^o, V_{it}^T < V_{it}^o, \Theta_1, \Theta_2, x_i, z_i) * \Pr(V_{i1}^T > V_{i1}^o,...,V_{it-1}^T > V_{it-1}^o, V_{it}^T < V_{it}^o \mid \Theta_1, \Theta_2, x_i, z_i).$$

The form of V_{it}^T will depend on the model that is being estimated, and increases in complexity as we go from the base to the learning to the forward-looking learning model. In addition, both x_{Si} and x_{oit} and ε_{Sit} and ε_{oit} enter linearly into V_{it}^T - V_{it}^o in equation (16), so that only contrasts between the two are identified. We assume that Var(ε_{Sit} - ε_{oit}) = 1 because the scale of the difference in values is not observed.

Directly maximizing (16) is difficult because the sequence of subject-period observations is not independent for individual i. Consequently, we adopt a two-step simulation approach that first estimates the choice probabilities given by the second term on the right-hand side of equation (16). Given a consistent estimate of the choice parameters Θ_1, the second step then estimates the health status equation accounting for dropout.

Estimation of dropout probabilities in the first step begins by noting that h_{it} is a function of Θ_2 and z_i. Consequently, the observed sequence of health status variables h_{it} is generated from

the true value of Θ_2. We can obtain estimates of the parameters of the dropout equation, Θ_1, conditioning on the observed values of h_{is}, by maximizing the log of

$$(17) \quad L_i^1 = \Pr(V_{i1}^T > V_{i1}^o, ..., V_{it-1}^T > V_{it-1}^o, V_{it}^T < V_{it}^o \mid \Theta_1, \{h_{is}\}, x_i)$$

using the GHK simulator because dropout is not independent across periods.

Given the estimated parameters Θ_1 from the first step, we now turn to the estimation of the density of $\{h_{it}\}$ in the second step. The joint density of the observed health status measures, conditional upon dropout in period t, is given by

$$(18) \quad L_i^2 = f(h_{i1}, ..., h_{it-1} \mid d_{i1}^T = 1, ..., d_{it-1}^T = 1, d_{it}^T = 0, \Theta_1, \Theta_2, x_i, z_i)$$
$$= \frac{f(h_{i1}, ..., h_{it-1} \mid \Theta_2, z_i)}{\Pr(d_{i1}^T = 1, ..., d_{it-1}^T = 1, d_{it}^T = 0 \mid \Theta_1, \Theta_2, x_i, z_i)}.$$

In order to recover estimates of Θ_2, we first substitute the estimated value of Θ_1 obtained from the first-step dropout equation into equation (18). We then estimate the log of equation (18) over all subjects to obtain the parameter estimates of the health status equation, using the GHK simulator to evaluate the joint probabilities.

5. RESULTS

We fit the two alternative models to the ACTG 175 data using the simulation methods described above. Health status at time t, H_{it}, is measured by the subject's CD4 count. The literature on ACTG 175, such as Hammer et al (1996), suggests that the following variables might be expected to affect the subject-level treatment impact on h_{it}, and hence are included in z_i: (1) demographic variables such as age, gender, and race; (2) variables measuring the extent of disease at baseline, such as whether the subject has a symptomatic HIV infection (which suggests a greater spread of the disease) and whether the subject had been exposed to prior antiretroviral therapy (which again may suggest a greater spread of disease).

A key feature of the empirical analysis is the specification of the factors that influence the

outside option available to trial participants at time t. Perhaps the most important factor is the set of AIDS treatments that become available over the course of the trial. While we do not observe such treatments directly, we assume that their availability is correlated with calendar time. The vector x_{iot} thus contains a variable indicating the current calendar year of the trial, as well as an indicator of whether the subject is an IV drug user, since these individuals are likely to be less stable and follow up in the trial (the other categories are haemophiliac and homosexual). It may be the case that subjects with more financial resources are more able to search and afford non-ACTG 175 treatments. While income is not reported, subject age, gender, and race may be correlated with the financial resources of the patient, and so we include these variables as proxies. Finally, we include the week number of the trial in x_{iot} because as the trial proceeds, subjects may be less likely to search for outside treatment options.

Side effects are a specified to be a function of treatment assignment (AZT+ddC vs. AZT), and as shown in equation (9) the impact of the side effects is assumed to be subject-specific. In addition, we allow the side effects associated with each treatment to differ depending on whether the subject has received antiretroviral therapy prior to ACTG 175 and whether the individual has a symptomatic HIV infection. We now turn to the discussion of the parameter estimates from the learning and forward-looking learning models, and then follow with a discussion of the treatment effects obtained from the models.

5.1 RESULTS FOR THE LEARNING MODEL

The parameter estimates from the learning model are presented in Table 3. The first column of the table presents coefficient estimates for the parameters that affect the decision to remain in ACTG 175 or drop out and seek alternative treatment. Aside from the side effects associated with the two treatments described below, the estimates from the dropout equation exhibit some notable findings. The positive estimate of γ implies that subjects with higher expected CD4 counts are more

likely to remain in the trial. In the absence of side effects, and with perfect information concerning the impact of treatment on CD4 counts, this would imply that subjects receiving combination therapy should be more likely to remain in the RCT than those receiving AZT. Moreover, declining CD4 counts over the course of the trial appear to drive attrition behaviour, since the positive coefficient on the week number of the trial suggests that individuals are more likely to remain the longer they have been in ACTG 175.

With regard to the coefficient estimates for the outside option parameters, intravenous drug users are much more likely to drop out of the trial in each treatment arm. If these individuals are to be studied further, incentives must be provided in order to induce them to remain in the trial. AIDS treatments available outside ACTG 175 appear to become increasingly attractive over time, since subjects who enrolled later in the trial are more likely to drop out. Finally, older individuals tend to be less likely to drop out of the trial, perhaps because they are more settled.

Another notable result in Table 3 concerns the estimates of the subjects' prior mean and variance of the effectiveness of trial treatment, μ_0 and $\sigma_{\mu 0}$. The positive and significant estimate of μ_0 in column (1) shows that the average subject appears to be optimistic about treatment effectiveness at the initiation of the trial. Moreover, the estimate of $\sigma_{\mu 0}$ suggests that most subjects expect that participation in the trial will have a positive impact on their health status. In order to assess the speed with which expected beliefs concerning the impact of the trial drugs on CD4 counts converge to the actual values, the expected ln(CD4) count in period in period, equal to $h_{i0} + \mu_i^t$, was constructed for representative subjects receiving combination therapy and AZT, respectively, assuming that each has a baseline CD4 count of 350.

Figure 5 shows that the expected ln(CD4) counts of both subjects in week 8 are equal and overstate the actual CD4 count, particularly for the AZT patient. This result is probably not

surprising, given that individuals choosing to enrol in the trial will likely have relatively optimistic expectations of effectiveness. The observed difference in CD4 counts in week 8 should therefore have little impact on the difference in dropout behaviour across treatment arms. As the subjects accumulate information on their CD4 counts over the course of the trial, their beliefs converge on the actual values, implying that the greater effectiveness of AZT+ddC eventually leads to reduced dropout compared to those receiving AZT alone. In addition, to the extent that subjects are initially over-optimistic about the likelihood of success, dropout will be relatively low during the early periods of the trial as subjects accumulate information concerning the effectiveness of treatment.

The posterior variance described in equation (12), $\sigma_{\mu,t}^2$, describes the degree of uncertainty regarding treatment effectiveness perceived by trial subjects. The relatively small value of σ_υ^2 shown in Table 3 suggests that the signals received by patients are relatively precise. Figure 6 describes the evolution of $\sigma_{\mu,t}^2$ over the course of the trial. The figure suggests that while subjects initially experience substantial uncertainty regarding effectiveness, by week 44 of the trial the posterior variance is roughly one-fifth the prior variance. By the end of the second year of the trial, $\sigma_{\mu,t}^2$ has dropped to approximately 0.02. Overall, Figures 5 and 6 suggest that while subjects are initially optimistic yet uncertain regarding the impact of the trial on health status, expected effectiveness converges on its observed value over the course of the trial, and patient uncertainty declines substantially.

5.2 RESULTS FOR THE FORWARD-LOOKING LEARNING MODEL

Not yet completed

5.3 TREATMENT EFFECTS

In this section we examine two sets of treatment effects, the first relating to the impact of AZT+ddC vs. AZT alone on health status, the second involving the impact of the alternative trial

treatments on side effects. Turning first to the impact of combination vs. monotherapy on CD4 counts, column (2) of Table 3 presents the estimates for the outcome equation (8). The results indicate that the average subject (who is not symptomatic nor has prior anti-retroviral therapy) receiving combination therapy experiences approximately a 5.4% increase in CD4 count in the first period (week 8) of the trial, compared with a 3.3% decline for those individuals receiving AZT alone. If all subjects were to remain in the trial for the full two years, the model estimates predicts that the treatment effect grows from 8.7% in week 8 to 31.9% by week 104. However, treatment effectiveness of combination therapy appears to be more uncertain than that of AZT alone, given the larger subject-level standard deviation associated with AZT+ddC. Demographic variables and symptomatic status appear to have little impact on CD4 counts, although individuals who received anti-retroviral therapy prior to ACTG 175 have lower CD4 counts regardless of trial arm. Finally, examination of columns (2) and (3) of Table 2 indicate that estimates of the CD4 equation obtained from the learning model tend to be similar to those obtained from a simple OLS regression, with the exception that CD4 counts among AZT patients are predicted to have a sharper decline when using the learning model.

To compare the impact of the alternative treatments on CD4 counts obtained from our model with those recovered from the raw data, we replaced the missing ln(CD4) counts of subjects who dropped out of the trial with values generated from the model in Table 3. Figure 7 presents the change in ln(CD4) count profiles after substituting for the missing values, along with the unadjusted profiles from Figure 1. Accounting for attrition has only a minor effect on patients receiving AZT+ddC. Despite the fact that half of these patients left by the end of the trial, the mean adjusted and unadjusted outcomes in week 104 are virtually identical. On the other hand, the plots for AZT patients show a marked difference between the adjusted and unadjusted profiles, suggesting that the unadjusted profile understates the disease progression for this group.

While the results indicate that combination therapy has a significant impact on CD4 counts, the picture appears to be quite different with regard to side effects. In this case, the results from column (1) of Table 3 suggest that patients receiving AZT+ddC are substantially more likely to drop out of the trial, after accounting for expected impact on CD4 counts, than are the AZT subjects, suggesting that side effects are greater for combination therapy. The box plots of the subject-level side effect distributions for combination therapy vs. monotherapy patients shown in Figure 8 confirm that most AZT patients are less likely to drop out of ACTG 175, all else equal, than the average AZT+ddC subject. One interpretation of our findings that is consistent with the difference in the treatment specific survivor functions plotted in Figure 2 is that AZT+ddC subjects are initially more likely to drop out of ACTG 175 due to the immediately perceived higher side effects associated with combination therapy. Over time, the fact that AZT subjects experience greater declines their CD4 counts offsets the lower side effects associated with the treatment, leading to increased attrition.

We have interpreted the difference in dropout probabilities across the treatment arms of ACTG 175 as reflecting the difference in side effects, conditional on expected CD4 count. Because side effects are not observed directly, one might argue that the results reflect differences in the outside treatment options available to subjects in the two groups. However, randomization implies that the average outside option will not differ across the two groups at baseline due to unobserved characteristics.

6. CONCLUSION

This paper assesses the impact of HIV combination therapy versus monotherapy for AIDS using data from the AIDS randomised clinical trial ACTG 175. We adopt a structural econometric framework that views subjects as utility maximizing agents who decide each period whether to drop out or remain in ACTG 175. Subject utility is specified to be a function of both long-term health

status, as measured by CD4 counts, and side effects. We consider models in which subjects learn over time about the impact of trial participation on CD4 counts, as well as allowing for forward-looking behaviour. The approach taken in this paper recovers treatment effect estimates for both CD4 counts and side effects. The framework also explicitly accounts for heterogeneous response to the treatment at the subject level in order to account for the possibility that patient sub-groups vary in responsiveness to treatment.

An examination of the ACTG 175 data indicates that subjects who leave the trial experience large declines in their CD4 counts prior to exit, particularly for those receiving AZT alone. Surprisingly, despite the fact that AZT+ddC appears to have a bigger impact on health status from the onset of the trial, combination therapy patients are initially slightly more likely to drop out. Eventually, however, attrition among AZT patients surpasses that of the AZT+ddC group. The (preliminary) estimates from the learning model indicate that AZT+ddC has greater side effects than AZT alone, implying higher attrition. Subjects appear to have optimistic prior beliefs concerning the effectiveness of trial treatment, so that initial differences in CD4 counts across treatments will have less impact on dropout. However, given that the estimated treatment effect of AZT+ddC increases from approximately 9% to 32% by the end of the trial, attrition among monotherapy patients overtakes that of combination therapy subjects due to the magnitude of the treatment effect on CD4 counts.

An important limitation of the models and estimates reported in this paper is the lack of information on the outside option available to subjects. Calendar time has a significant impact on the dropout decision, suggesting that alternative therapies become more attractive over the course of the trial period. The results are conditional upon the decision of subjects to participate in ACTG 175, and cannot be generalized to non-participants since their outside option (or beliefs about the effectiveness of trial treatment) is likely to differ from those of trial patients. If data on non-

participants could be obtained, the decision to enrol in the trial could be incorporated into the framework, and perhaps the findings of treatment effectiveness could be generalized to the HIV population (with CD4 counts between 200 and 500) as a whole.

TABLE 1
OLS ESTIMATES OF THE IMPACT OF AZT+ddC ON CD4 COUNTS
Dependent Variable is Change in ln(CD4) Count Relative to Baseline

Variables	Specification		
	(1)	(2)	(3)
AZT+ddC, Week 8	0.100	0.101	0.089
	(0.019)	(0.019)	(0.018)
AZT+ddC, Week 20	0.126	0.130	0.115
	(0.020)	(0.020)	(0.020)
AZT+ddC, Week 32	0.127	0.130	0.116
	(0.021)	(0.021)	(0.021)
AZT+ddC, Week 44	0.140	0.143	0.129
	(0.022)	(0.022)	(0.021)
AZT+ddC, Week 56	0.152	0.157	0.144
	(0.023)	(0.023)	(0.022)
AZT+ddC, Week 68	0.166	0.170	0.159
	(0.024)	(0.024)	(0.024)
AZT+ddC, Week 80	0.217	0.222	0.216
	(0.026)	(0.025)	(0.025)
AZT+ddC, Week 92	0.177	0.179	0.178
	(0.027)	(0.026)	(0.026)
AZT+ddC, Week 104	0.191	0.192	0.191
	(0.028)	(0.027)	(0.026)
Week 8	-0.032	-0.029	-0.026
	(0.020)	(0.020)	(0.020)
Week 20	-0.070	-0.067	-0.067
	(0.021)	(0.020)	(0.021)
Week 32	-0.114	-0.112	-0.120
	(0.021)	(0.021)	(0.021)
Week 44	-0.138	-0.137	-0.157
	(0.022)	(0.021)	(0.022)
Week 56	-0.140	-0.142	-0.173
	(0.022)	(0.022)	(0.022)
Week 68	-0.148	-0.150	-0.192
	(0.023)	(0.023)	(0.023)
Week 80	-0.174	-0.177	-0.232
	(0.024)	(0.023)	(0.024)
Week 92	-0.156	-0.159	-0.225
	(0.025)	(0.024)	(0.025)
Week 104	-0.158	-0.160	-0.235
	(0.025)	(0.025)	(0.025)
Includes Covariate Controls?	No	Yes	Yes
Includes Indicators for Week of Dropout?	No	No	Yes

Notes: Standard errors in parentheses. Based on 7063 subject-week observations. Covariate controls include age, gender, race, symptomatic HIV infection, screening CD4 count, and prior antiretroviral therapy.

TABLE 2
REPORTED REASON FOR DROPOUT, CONDITIONAL UPON ATTRITION

Reason	AZT+ddC			AZT		
	Overall	Dropout in Weeks 8-44	Dropout in Weeks 56+	Overall	Dropout in Weeks 8-44	Dropout in Weeks 56+
	(1)	(2)	(3)	(4)	(5)	(6)
Death	0.02	0.01	0.04	0.04	0.04	0.03
Toxicity of Treatment, Patient Request	0.27	0.34	0.13	0.19	0.22	0.16
Request of Patient	0.26	0.23	0.31	0.33	0.32	0.34
Request of Investigator	0.03	0.04	0.02	0.02	0.03	0.01
Lost to Follow-Up	0.15	0.16	0.13	0.12	0.13	0.10
Other	0.27	0.22	0.37	0.30	0.26	0.36

TABLE 3
PARAMETER ESTIMATES FOR THE LEARNING MODEL

Variables	Dropout Equation (1)	$\ln(CD4_{it}) - \ln(CD4_{i0})$ (2)	OLS Estimates (3)
γ	0.203		
	(0.030)		
AZT (mean)	-0.323	-0.005	0.030
	(0.159)	(0.001)	(0.041)
AZT+ddC (mean)	-0.696	0.053	0.114
	(0.154)	(0.0008)	(0.038)
AZT (s.d.)	0.211	0.078	
	(0.069)	(0.031)	
AZT+ddC (s.d.)	0.168	0.129	
	(0.055)	(0.052)	
AZT*t		-0.028	-0.023
		(0.002)	(0.002)
AZT+ddC*t		0.001	-0.004
		(0.0014)	(0.004)
Prior Antiretroviral	0.147	-0.077	-0.099
*AZT	(0.017)	(0.022)	(0.047)
Prior Antiretroviral	0.114	-0.129	-0.116
*AZT+ddC	(0.014)	(0.018)	(0.039)
Symptomatic*AZT	-0.109	-0.056	-0.030
	(0.023)	(0.182)	(0.116)
Symptomatic	-0.146	0.011	-0.019
*AZT+ddC	(0.017)	(0.087)	(0.212)
Age	0.016	-0.0001	0.0003
	(0.001)	(0.007)	(0.044)
Male	-0.179	0.022	0.0009
	(0.028)	(0.086)	(0.062)
White	0.054	0.002	-0.009
	(0.020)	(0.017)	(0.013)
IV Drug User	-0.545		
	(0.290)		
Homosexual	0.096		
	(0.026)		
t	0.399		
	(0.024)		
Calendar Year	-0.387		
	(0.019)		
μ_0	0.476		
	(0.144)		
$\sigma_{\mu 0}$	0.243		
	(0.045)		
σ_{υ}		0.131	
		(0.002)	

Note: Standard errors in parentheses. Estimates based on 536 subjects receiving AZT and 536 subjects receiving AZT+ddC.

REFERENCES

Boscardin, W., Taylor, J., and Law, N. (1998), "Longitudinal Models for AIDS Marker Data," *Statistical Methods in Medical Research*, 7, 13-27.

Ching, A. (2000), "Dynamic Equilibrium in the U.S. Prescription Drug Market After Patent Expiration," mimeo, Department of Economics, University of Minnesota.

Crawford, G. and Shum, M. (2000), "Uncertainty and Learning in Pharmaceutical Demand," mimeo, Dept. of Economics, Duke University.

Currie, G. and Park, S. (2000), "The Effects of Advertising and Consumption Experience on the Demand for Antidepressant Drugs," mimeo, Dept. of Economics, University of Calgary.

Erdem, T. and Keane, M. (1996), "Decision-Making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets," *Marketing Science*, 15, 1-20.

Hammer, S., et. al. (1996), "A Trial Comparing Nucleoside Monotherapy with Combination Therapy in HIV-Infected Adults with CD4 Cell Counts from 200 to 500 per Cubic Millimeter," *New England Journal of Medicine*, 335, 1081-1090.

Heckman, J. and Robb, R. (1985), "Alternative Models for Evaluating the Impact of Interventions," in *Longitudinal Analysis of Labor Market Data*, eds. J. Heckman and B. Singer, Cambridge University Press, New York, 156-245.

Philipson, T. and Desimone, J. (1997), "Experiments and Subject Sampling," *Biometrika*, 84, 619-631.

Philipson, T. and Hedges, L. (1998), "Subject Evaluation in Social Experiments," *Econometrica*, 66, 381-408.

Rubin, D. (1974). "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies," *Journal of Educational Psychology*, 66, 688-701.

Rubin, D. (1978), "Bayesian Inference for Causal Effects," *The Annals of Statistics*, 6, 34-58.

Scharfstein, D., Rotnitzky, A., and Robins, J. (1999), "Adjusting for Non-Ignorable Drop-Out Using Semiparametric Nonresponse Models (with discussion)," *Journal of the American Statistical Association*, 94, 1096-1146.

FIGURE 1
CD4 COUNT PROFILES, BY TREATMENT GROUP

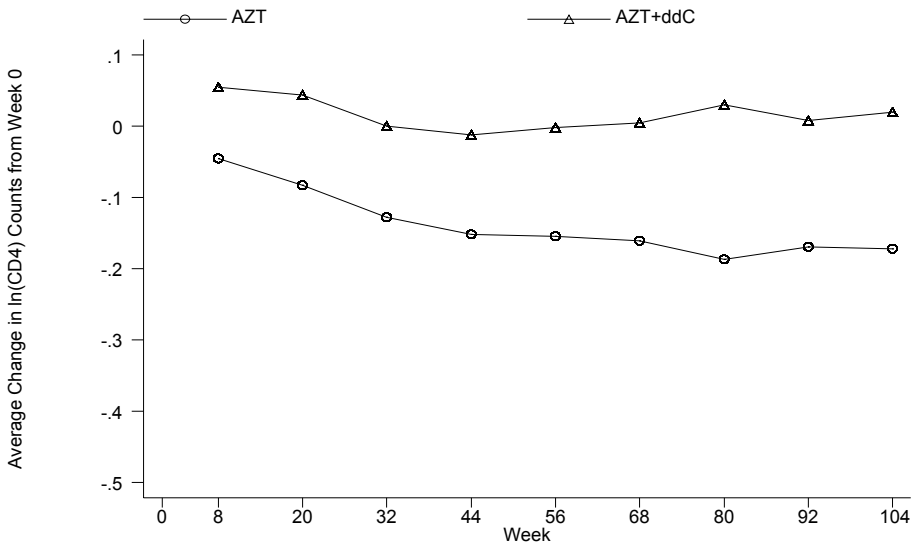

FIGURE 2
SURVIVOR FUNCTIONS, BY TREATMENT GROUP

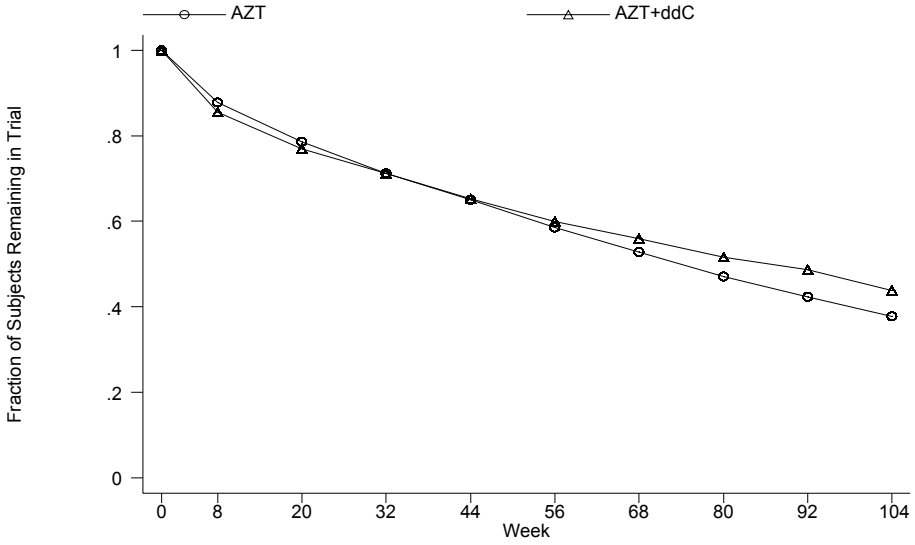

FIGURE 3
CD4 COUNT PROFILES FOR AZT+ddC SUBJECTS, BY ATTRITION GROUP

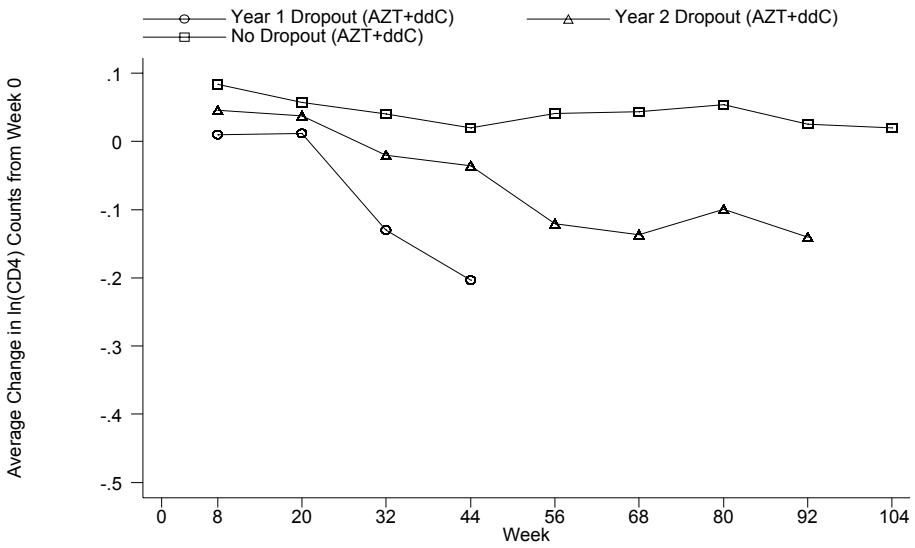

FIGURE 4
CD4 COUNT PROFILES FOR AZT SUBJECTS, BY ATTRITION GROUP

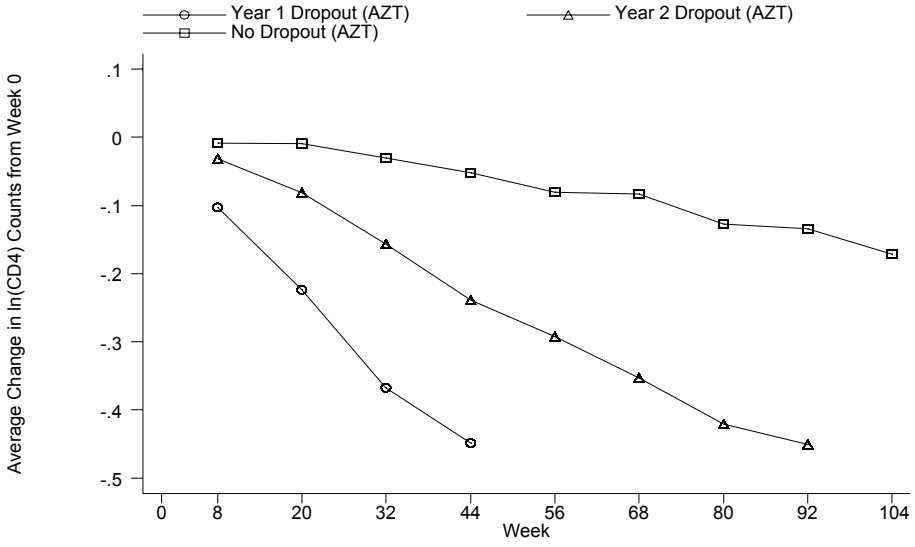

FIGURE 5
ln(CD4) COUNT PROFILES, ACTUAL AND EXPECTED, BY TREATMENT GROUP

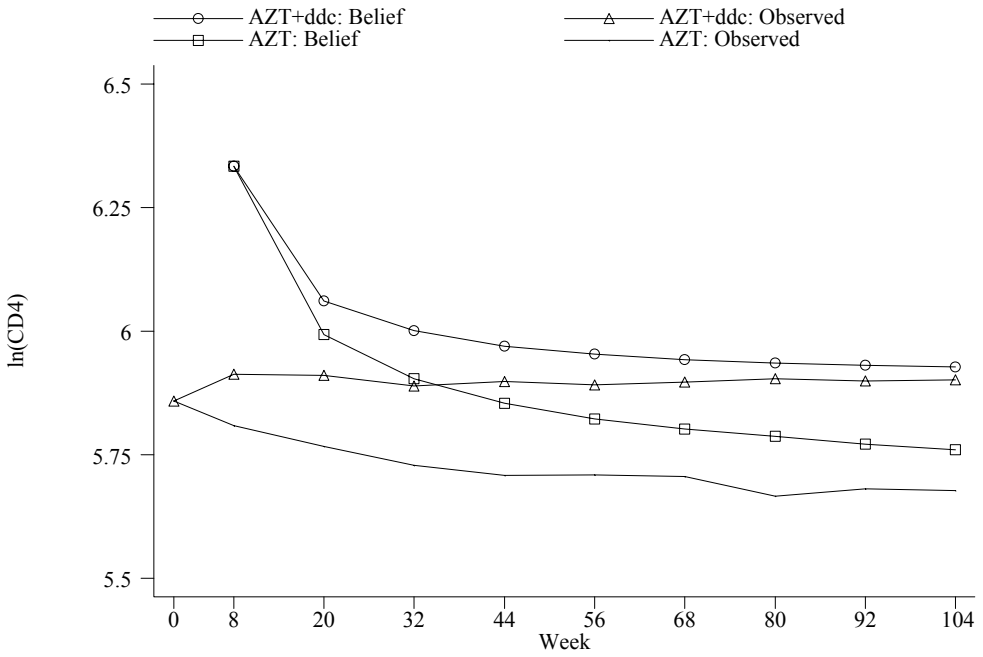

FIGURE 6
EVOLUTION OF POSTERIOR VARIANCE

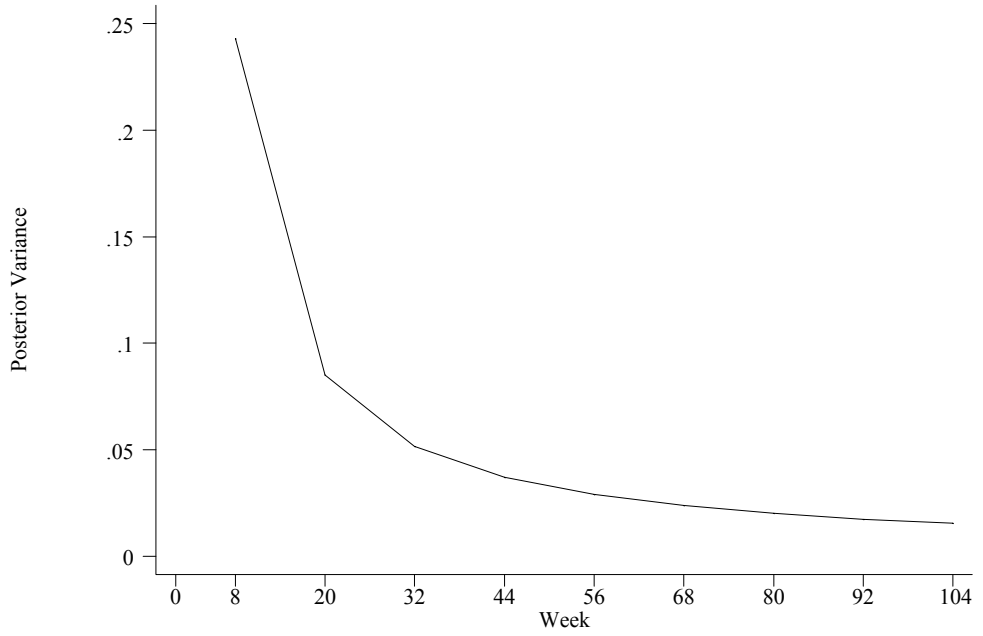

FIGURE 7
CD4 COUNT PROFILES, ADJUSTED FOR ATTRITION, BY TREATMENT GROUP

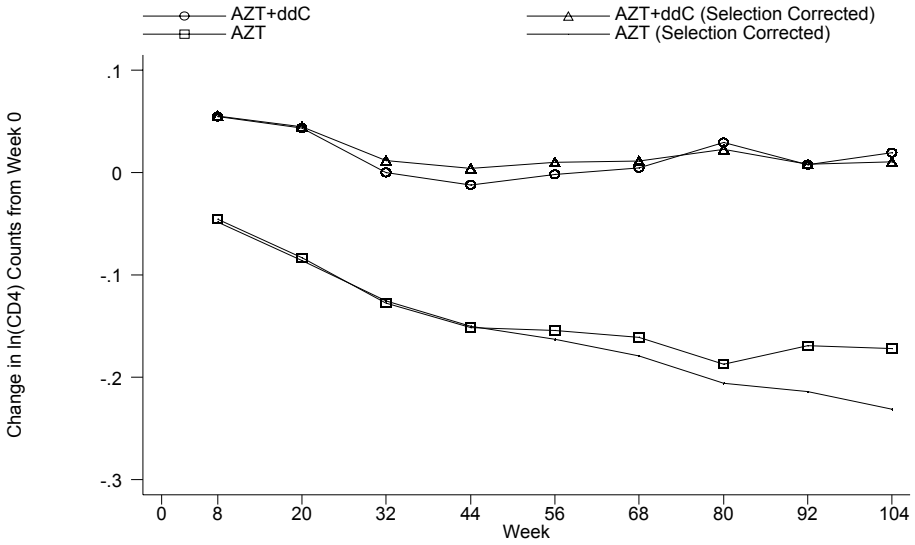

FIGURE 8
DISTRIBUTION OF SUBJECT SIDE EFFECTS, BY TREATMENT GROUP

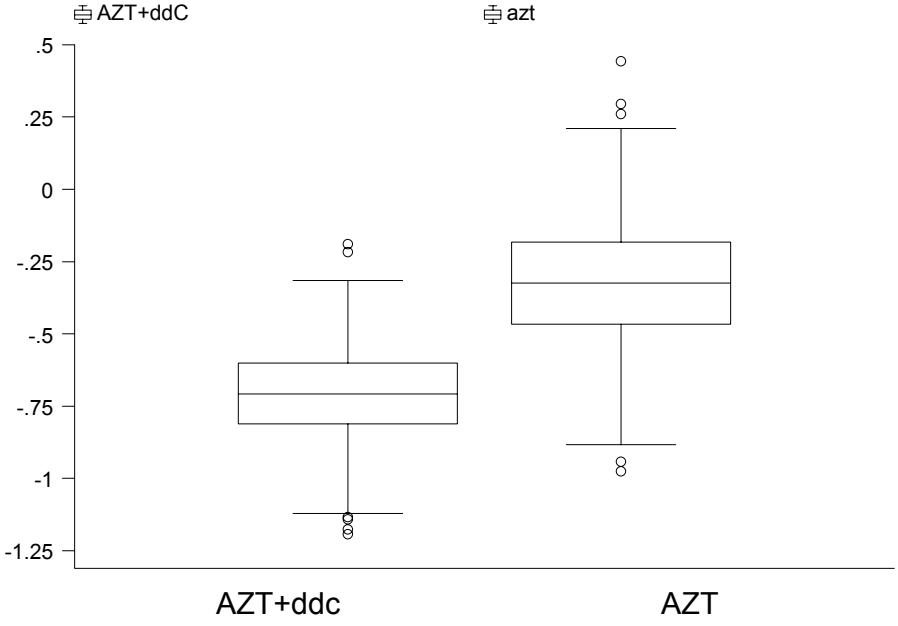

APPENDIX TABLE 1
SUMMARY STATISTICS BY TREATMENT STATUS

Variable	AZT	AZT+ddC
Age at Baseline	34.54 (8.75)	34.78 (8.65)
Male	0.80 (0.40)	0.84 (0.37)
White	0.70 (0.46)	0.69 (0.46)
Symptomatic HIV Infection	0.17 (0.38)	0.18 (0.38)
Screening CD4 Count	348.97 (83.78)	350.04 (82.34)
Prior Antiretroviral Therapy	0.55 (0.50)	0.55 (0.50)
IV Drug User at Baseline	0.12 (0.33)	0.14 (0.35)
Homosexual at Baseline	0.62 (0.49)	0.67 (0.47)
Number of Observations	536	536
Number of Sites	89	87

Note: Standard deviation in parentheses.